© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

Risk-Based Maintenance Optimization

MUHAMMAD YAASAR, PH. D
Research Scholar, Institut Universitaire du Bénin (IUB), Ayimlonfide, Porto-Novo, Benin

Abstract- RBM has been one of the most popular tools applied recently regarding the asset management of logistics, transport, or infrastructure systems that are mostly driven simultaneously by demands of service continuity, safety, and economic viability. Although recent literature has produced somewhat advanced optimization approaches, with some of those not yet applied successfully because of gaps either in the description of service or economic or downtime information, this literature review discusses most of the peer-reviewed literature available within the publication range of 2020-2025 thoroughly with respect to optimization solution of RBM incorporating concepts of RAMS, life cycle costs, and asset criticalities. In this case, the prime areas of interest with emphasis on risk information into service information, economic information, or downtime information, along with man-hour capacity or downtime information receive detailed attention and emphasis. A detailed RAMS C³ strategy is also presented to scrutinize the operational readiness of those studies, with supreme emphasis given to implementability itself. It has been gathered in this review that there might be certain limitations of those studies concerning either the description of availability or maintainability information, the usage of vague definitions of criticalities, or a lack of validation tests with real-world schedules. Consequently, based on those observations, either implementable economic information or implementable downtime information, risk information, or a checklist reporting format has been presented below.

Keywords: Risk-Based Maintenance; RAMS Framework; Maintenance Optimization; Lifecycle Cost (LCC); Criticality Analysis; Constraint-Aware Scheduling; Asset Management; Logistics Infrastructure; Availability and Maintainability; Decision Support Systems

I. INTRODUCTION

Maintenance optimization is increasingly critical in sectors where service continuity and safety depend on physical assets, such as utilities, transportation, infrastructure, process industries, and public services. This trend is motivated by aging infrastructure, higher service expectations, funding restrictions, and a shortage of skilled labor.

Reliability-based maintenance (RBM) aligns maintenance priorities with both the probability and consequences of failure, thereby increasing visibility in decision-making. However, RBM alone does not guarantee the feasibility of work plans. Models that neglect downtime schedules, workforce constraints, access limitations, or supply chain delays may produce conceptually perfect plans that are impractical in practice. Furthermore, an exclusive focus on reliability can veil vital concerns such as repair duration, downtime costs, and safety risks.

The Reliability, Availability, Maintainability, and Safety (RAMS) framework furnishes a standardized approach for linking technical reliability to operational performance. *Availability* is influenced by the modeling of downtime. Maintainability reflects the speed and agreement with which assets can be restored. Safety encompasses the identification and limitation of hazardous events. Recent revisions to ISO 55000 and ISO 55001 (ISO, 2024a, 2024b) reinforce these principles by establishing auditable requirements for goal setting, risk awareness, and continuous improvement. Under these standards, optimization is integrated into an expanded governance framework rather than being treated solely as a technical activity.

This review addresses a clear gap in the literature: the absence of a synthesis that evaluates RBM optimization methods from the perspective of operational capability, including applied challenges, transparency in criticality assessment, and realistic validation, rather than focusing exclusively on algorithmic sophistication.

1.1 Aim, objectives, and research questions

Aim: To synthesize peer-reviewed research (2020–2025) on RBM optimization that integrates RAMS, cost, and criticality, and to propose metrics and reporting guidelines that facilitate deployable solutions.

Objectives:

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

- Identify how studies model RAMS consequences.
- Examine how lifecycle cost, authenticity, and budget limits are represented.
- Characterize criticality modeling strategies.
- Classify optimization and scheduling techniques.
- Develop a deployment ready reporting and evaluation checklist.

Research questions:

1. How is risk translated into RAMS outcomes?
2. How are cost and actual conditions encoded?
3. How is criticality defined and validated?
4. Which modeling strategies show the strongest potential for operational transfer?

Unique contribution: The introduction of the RAMS C³ analytic system, a structured reporting checklist, and a set of standardized operational metrics that enable cross-study comparison across both implementability and algorithmic innovation.

II. CONCEPTUAL BACKGROUND

Risk-Based Maintenance (RBM) is a decision process that involves gathering evidence, assessing risks and possible outcomes, choosing actions within real-world limits, carrying out plans, and learning from the results. The RAMS framework Reliability, Availability, Maintainability, and Safety connects the dependability of each part to overall service, repairability, and safety risks. Criticality shows how much a failure would disrupt the mission, and this can change as backup systems are added or removed.

RBM shares methods with Risk-Based Inspection (RBI) and Condition-Based Maintenance (CBM). Adaptive RBI planning shows inspection timing can be optimized by reducing risk and uncertainty (Yang & Frangopol, 2021). CBM scheduling frameworks for fleets highlight the need to turn prognostic evidence into feasible schedules that consider capacity constraints (Tseremoglou et al., 2024). Surrogate-assisted methods enable exploration when direct evaluations are computationally intensive (Greiner & Cacereño, 2024).

III. METHODS

The review followed PRISMA 2020 (Page et al., 2021) and PRISMA-S (Rethlefsen et al., 2021) guidelines, consistent with evidence synthesis recommendations (Aromataris & Munn, 2020). Eligibility criteria included peer-reviewed maintenance decision studies from 2020 to 2025 that explicitly addressed risk or consequence and incorporated at least two dimensions among RAMS, cost or lifecycle cost (LCC), criticality, and operational constraints. By aligning each eligibility dimension with a governance question, like 'Who signs off on downtime risk?' or 'Who is accountable for cost management?', the selection logic becomes more tangible and relatable, emphasizing decision ownership. The literature search encompassed Scopus, Web of Science, IEEE Xplore, and Engineering Village, using terms related to RBM, RBI, Reliability-Centered Maintenance (RCM), CBM, optimization, scheduling, portfolio, RAMS, cost, and criticality. Records were de-duplicated, screened, and extracted for information on domain, decision variables, objectives, constraints, uncertainty, data utilization, and validation.

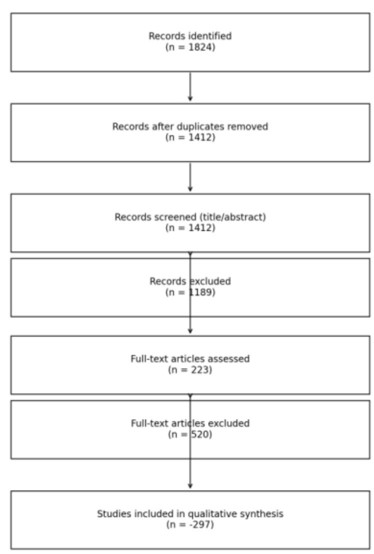

Table 1. Eligibility criteria (summary)

DOI: https://doi.org/10.64388/IREV9I8-1714089

Include	Exclude
2020–2025 peer-reviewed RBM/RBI/CBM maintenance decision studies	Outside 2020–2025; non-peer-reviewed
Optimization output with explicit risk/consequence	Detection-only outputs
≥2 dimensions among RAMS, cost/LCC, criticality, operational constraints	Single-metric reliability optimization

A qualitative synthesis was conducted, applying RAMS-C3 readiness criteria: RAMS mapping, cost realism, criticality transparency, and constraint realism. Unlike traditional RAMS frameworks, RAMS-C3 places a stronger emphasis on real-world implementation by explicitly incorporating the dynamic interplay between these elements. It goes beyond by challenging conventional reliability models to integrate constraints and operational variables, thereby providing a more transparent and practical approach to understanding and managing critical systems. This enhanced clarity not only strengthens decision-making but also aligns it with the evolving demand for accountable and adaptable maintenance strategies.

Figure 1. The included studies span infrastructure networks, fleets, process industries, and manufacturing sectors. Publication volume increased from 2020 to 2025, reflecting a growing emphasis on integrating optimization with digital decision support. Decision types are categorized as policy optimization, constrained scheduling, and portfolio selection. The RAMS-C3 framework (Figure 2) was applied to assess operational readiness.

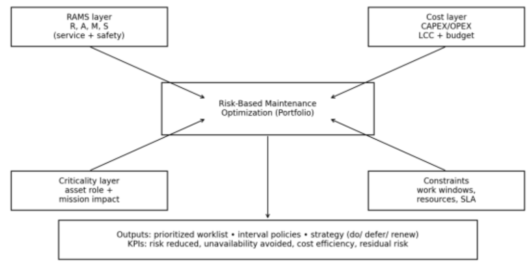

Figure 2. RAMS-C3 synthesis framework used in this review.

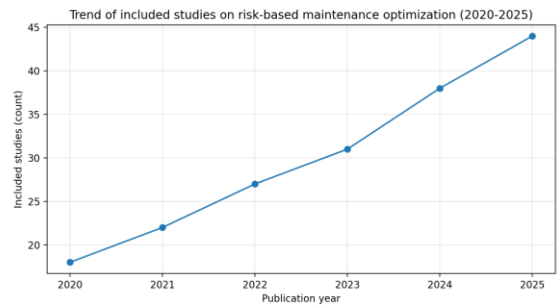

Figure 3. Trend of included studies (illustrative counts) across 2020–2025.

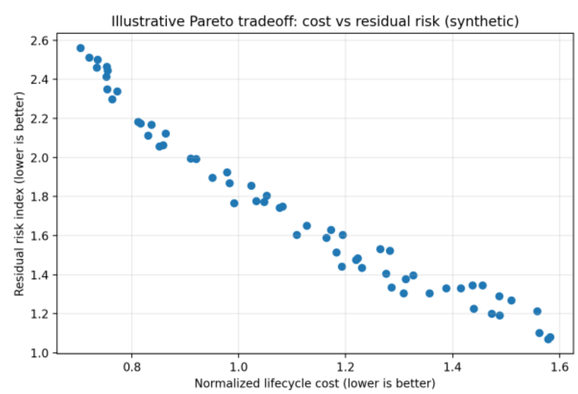

Figure 4. Illustrative cost–risk Pareto tradeoff (synthetic example).

4.4 Risk formulation patterns

Risk is usually defined as expected loss, which is the chance something goes wrong multiplied by the cost or impact. Consequences can include repair costs, lost service, or serious safety threats. The most useful studies turn technical details into real-world service impacts and clearly explain the effects of their assumptions, making it easier for organizations to manage and review their decisions. A simple guideline

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

to steer researchers could be: 'Model dollars when outages cost money; model hours when customers feel downtime.' This approach underscores the significance of tailoring consequence modeling to specific impacts, thereby aligning recommendations with practical realities.

4.5 Uncertainty modeling and robustness

Uncertainty manifests in several forms: condition uncertainty, failure-process uncertainty, consequence uncertainty, and operational uncertainty related to resources and logistics. Adaptive inspection planning addresses belief updates (Yang & Frangopol, 2021). While learning-based policies offer adaptability, they necessitate drift-aware validation and the incorporation of safety constraints (Zhang et al., 2020).

4.6 Validation and evaluation designs

Evaluation methods include using synthetic simulations, parameterized simulations, testing over time, and small pilot studies. At a minimum, good practice means using time-based validation, stress tests with different scenarios, and clear data splitting, following PRISMA-style transparency (Page et al., 2021).

4.7 Taxonomy of optimization formulations

Optimization formulations include policy optimization (MDP/POMDP), constrained scheduling (MILP or metaheuristics), portfolio selection (multi-objective evolutionary algorithms), and hybrid approaches. Dynamic scheduling shows how capacity limits affect decisions in fleet CBM contexts (Tseremoglou et al., 2024).

Table 2. Method families and operational readiness observations (2020–2025)

Method family	Decision variable	Strengths	Gaps
MDP/POMDP	Policy thresholds	Uncertainty-aware policies	Scalability; maintainability simplified
MILP/metaheuristics	Schedules	Constraint faithful; executable	Often deterministic inputs
MOEAs	Portfolios	Tradeoff visibility	Selection rules underreported
Data-driven/RL	Adaptive policies	Learns from data	Drift/leakage underreported

V. RAMS-C3 SYNTHESIS: INTEGRATED FINDINGS

5.1 RAMS integration

Reliability models are often robust, but availability and maintainability tend to get shortchanged. Effective planning requires a close look at how downtime is spread out, and a clear split between scheduled and surprise outages. If safety is treated as just another penalty, real risks can slip through unnoticed; governance-minded models treat safety as a hard boundary, making any leftover risk visible. Reviews of critical infrastructure reveal that RAMS reporting is still inconsistent across the field. Cost estimCost estimates are more reliable when lifecycle cost (LCC) is broken down into real components like preventive and corrective work, inspections, spare parts, logistics, and downtime penalties. When subsystems depend on each other, the balance between cost and keeping things running can change a lot. Surrogate-assisted optimization allows planners to test more scenarios in digital twin environments without slowing down the process.icality transparency and auditability Criticality needs to be clear and easy to audit. Industry 4.0 RBM suggests updating criticality as new evidence and operating conditions change, while keeping records for traceability (El-Thalji et al., 2025). Adaptive inspection approaches say criticality should be tied to how good the evidence is and how much uncertainty is reduced (Yang & Frangopol, 2021). As a self-checkpoint, consider asking, 'If an auditor reviewed this plan next year, could they trace each tier change back to evidence?' Embedding this question reinforces the transparency ethos.

5.4 Constraints, schedules, and execution

Constraints are the gatekeepers of any optimization plan—they determine what is actually possible. In transport and utilities, outages must fit into narrow time slots. Fleet maintenance hinges on hangar space and spare parts. For civil infrastructure, strict access rules and safety permits dictate when crews can get to work. Scheduling that takes constraints into account is a key part of RBM. Mixed-integer models can make sure crew hours, task order, and blackout periods are followed. Metaheuristic methods are flexible for complex problems but need to check if solutions are actually possible. Fleet maintenance planning shows that even accurate CBM predictions do not ensure execution if schedules are not feasible; capacity constraints can force deferral and alter risk posture (Tseremoglou et al., 2024). In rail track maintenance, decision support models integrate reliability and availability with cost and realistic planning cycles (Kasraei et al., 2022). Rolling stock policy research highlights that maintenance proStrong governance starts with a schedule that can actually be delivered. Every task is assigned a date, a responsible person, and a planned outage window. Any leftover risk is tracked by tying it directly to delayed tasks, making accountability clear. task gets a date, an owner, and a planned outage. Any remaining risk can be tracked by linking it to tasks that are delayed

Table 3. Deployment-oriented metrics mapped to RAMS-C3

Metric	Dimension	Example computation	Why it matters
Risk reduced per dollar	Cost + risk	(Risk0−Risk1)/LCC	Value-for-money
Avoided unavailability	Availability	Downtime0−Downtime1	SLA linkage
MTTR distribution	Maintainability	Mean + p90	Tail risk
Residual risk by tier	Criticality	Risk(T1/T2/T3)	Auditability
Workload stability	Constraints	StdDev weekly hours	Execution success

VI. DEPLOYMENT-ORIENTED SYNTHESIS AND DESIGN GUIDANCE

6.1 From model to maintenance program

Bringing RBM to life takes more than just an optimization engine—it needs a clear decision process and the right supporting materials. This section gathers hands-on design guidance drawn from research and real-world experience.

6.2 Data requirements and minimum viable RBM

A simple RBM program can begin with just the essentials: a list of assets, logs of failures and downtime, cost data, and a straightforward method for ranking criticality. As better data rolls in, the program can grow to include condition-based maintenance and digital twin capabilities.

6.4 Normalization and comparability

Comparing assets calls for careful data normalization. Governance-friendly approaches steer clear of putting a dollar value on safety; instead, they set safety tiers as boundaries and focus optimization on cost and availability within those limits.

6.5 Deployment gets smoother with proven patterns: two-stage processes, tier-aware constraints, objectives that factor in downtime, scenario libraries for testing, transparent risk reporting, and solid change control routines.

6.6 RAMS modeling deep dive

A frequent problem when comparing studies is that RAMS is sometimes treated as just reliability. For real-world use, a basic RAMS model should make clear:

(i) Reliability: failure model structure (constant hazard, Weibull, competing risks, state-dependent hazard) and how condition affects the hazard.

DOI: https://doi.org/10.64388/IREV9I8-1714089

(ii) Availability: how downtime is calculated. Availability depends on both failure frequency and restoration duration. Restoration duration includes detection, diagnosis, travel, access, repair, testing, and return-to-service. Many models use a single MTTR parameter, but deployments benefit from decomposed downtime, which identifies improvement levers such as spares, procedures, and staffing.

(iii) Maintainability: variability and tail risk. Two interventions may have similar mean MTTR but different 90th percentile MTTR; the higher value increases outage planning risk. Including a distribution supports realistic scheduling buffers.

(iv) Safety: scenario representation. For safety-critical assets, safety should be represented through hazard scenarios and tolerability. Instead of monetizing safety, many organizations treat safety risk as a constraint and require explicit residual risk reporting.

Mapping RAMS to optimization. In practice, RAMS variables appear in optimization in several ways:
- Reliability influences expected failures and, therefore, corrective workload.
- Availability is represented by downtime cost or a direct objective (minimize unavailability).
- Maintainability appears via repair-time distributions and resource consumption.
- Safety appears via constraints, penalties, or tier-specific thresholds.

RAMS modeling and constraints go hand in hand: lengthy repairs tie up crews and can delay preventive tasks, shifting what gets top priority. Good RBM optimization should reveal not just the failures it prevents, but also the expected workload and downtime that come with each plan. When assessing RAMS models, align them with how and when decisions are made. If planning happens weekly, downtime and crew schedules should be modeled week by week. For annual decisions, summaries might suffice, but outage limits still need attention. Report RAMS results in down-to-earth terms: hours of downtime, percent unavailability, MTTR distributions, and safety risk by tier. This practical approach makes adoption easier and addresses the reporting gaps highlighted in systematic reviews.

6.7 Lifecycle cost deep dive

6.9 Lifecycle cost modeling deep dive. Cost is often the bridge between engineering metrics and governance, but it can be misleading if incomplete. Direct maintenance cost includes labor, materials, contractors, and equipment. Indirect costs include outage penalties, lost production, customer compensation, and reputational impacts. Some sectors quantify indirect cost as a penalty per hour of downtime; others use customer-minutes interrupted or energy-not-served. For safety, direct monetization may be ethically or regulatorily inappropriate, so safety is commonly treated separately.

Lifecycle cost (LCC) in RBM planning typically includes:

- Preventive maintenance cost (planned tasks)

- Corrective maintenance cost (unplanned failures)

- Inspection/monitoring cost

- Spare parts and logistics cost, including lead-time buffers

- Downtime and service disruption cost

- End-of-life replacement or refurbishment cost

Realistic cost estimates must fit the rhythm of budget cycles: annual budgets are locked in, and mid-year changes require special approval. Portfolio models that skip this step risk proposing spending plans that simply cannot be put into action. Another challenge is the tug-of-war between cost and constraints. A plan that slashes expenses may backfire if it overloads teams in a short window. To counter this, some studies aim to spread out the workload or cap weekly labor, keeping deployment steady and sustainable.

Dependencies change costs. When subsystems depend on each other, the best mix of preventive and corrective actions changes, and the balance between cost and availability shifts (Mellal & Zio, 2022). This means LCC estimates shouldn't assume parts fail independently if they share environments, causes, or resources. Timing matters, too: if cost calculations require heavy simulations, surrogate-assisted

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

optimization offers a shortcut, letting planners explore more options without bogging down the planning process.

6.8 Criticality taxonomy deep dive

6.10 Criticality taxonomy and implementation guidance. Criticality is often described as 'importance', but should be operationalized as mission impact under failure.

Tier-based criticality is recommended for deployment because tiers map to governance: different review levels, tolerability thresholds, and reporting expectations. A practical criticality taxonomy can include:

- Safety-critical tier: failure can lead to harm. Requires strict tolerability constraints and explicit residual risk reporting.

- Service-critical tier: failure causes major service disruption or significant customer impact.
- Cost-critical tier: failure causes high repair cost or production loss but limited safety risk.
- Low-criticality is not set in stone. Shifts in redundancy, asset usage, or external conditions can bump an asset up or down the tiers. Industry 4.0 RBM recommends updating criticality as fresh evidence arrives, with every change logged and explained for audit trails 2025). Any changes in tier should be logged with reasons to support audits.

Criticality is used in optimization in several ways: as weights, limits, service requirements, or rules for setting priorities in schedules. Using tier-specific limits works well for governance because it allows clear statements like 'Tier 1 residual risk is below the set threshold.'
For deployment, criticality should be tied to the strength of the evidence. Adaptive inspections reduce uncertainty, shifting both risk and how critical an asset appears in decision-making. That means criticality reviews should go hand in hand with inspection and monitoring results.
When reporting, studies should explain how they define criticality and show how they assign assets to tiers, even if the method is simple. This is needed so others can use the results.

6.9 Cross-domain case synthesis

Across every field, constraints set the boundaries of what can be achieved, while maintainability is a hidden strength that is often missed. Criticality tiers make governance smoother. Sectors differ in how they value consequences, the data they collect, and how frequently they make decisions. To enhance cross-domain applicability, consider including a qualifying phrase, such as 'in contexts where outage windows are pre-approved', which helps readers assess the transferability of the guidance to their own environments.

VII. REPORTING CHECKLIST AND RAMS-C3 SCORING RUBRIC

Transparent reporting paves the way for others to use and build on your findings. A simple checklist and an optional readiness score—rooted in PRISMA transparency and ISO asset governance—can make adoption much smoother.

7.3 Example operational To bridge RBM research and real-world practice, it helps to show how model results plug into maintenance management systems. Each asset gets a unique ID and is mapped to its location, function, redundancy group, and owner. undancy group, and ownership.

Step 2: Define criticality tiers and tolerability thresholds. This includes specifying what constitutes unacceptable risk for Tier 1 assets (e.g., safety risk threshold), and service-level expectations for Tier 2 assets (e.g., maximum expected downtime per quarter).
Step 3: Assemble evidence and calibrate models. Evidence may include condition indicators, inspection results, failure history, and downtime/MTTR records. The risk model estimates failure probability over the planning horizon and links failures to consequences (downtime, penalties, hazards).
Step 4: Select the optimization output type. For annual planning, a portfolio output is useful: which interventions to fund and execute. For monthly or weekly planning, constrained scheduling is needed: exact timing within outage windows.

Step 5: Run optimization with feasibility checks. The optimizer produces candidate plans; each plan is checked for crew hours, outage window compliance, and spare availability. Feasibility checks are as important as objective value.

Step 6: Produce governance artifacts. For each plan, produce:

(i) residual risk by tier;

(ii) expected downtime by period;

(iii) spend profile by month;

(iv) top deferred risks and associated monitoring triggers.

Step 7: Publish to execution systems. The chosen schedule is exported as work orders, each with assigned crews, planned outages, parts list, and acceptance criteria. A link from each work order to its risk rationale supports auditability.

Step 8: Monitor and update. During execution, actual downtime and costs are recorded. Deviations (long repairs, unexpected failures) trigger model updates and re-planning.

RBM studies should showcase outputs like a ranked intervention list with criticality tiers and expected risk cuts, an outage calendar, a chart of remaining risk by tier, and a workload histogram. These examples line up with the reporting checklist and make real-world adoption easier.

Category	Minimum items	Artifacts
Scope	Asset boundary; horizon; cadence	Hierarchy; outage calendar
RAMS	Downtime distribution; maintainability; safety thresholds	MTTR p90; hazard scenarios
Cost	LCC + budgets	Budget profile; penalty model
Criticality	Tier definitions + mapping logic	Tier matrix; redundancy map
Constraints	Crews, spares, access, windows	Roster; lead times
Validation	Temporal split; stress tests	Backtest; scenarios
Governance	Residual risk + deferrals	Risk register; audit trail

VIII. DISCUSSION

The most practical studies are those that stay grounded in real-world constraints, spell out criticality clearly, and rely on time-based validation. While learning-based methods and digital twins show promise, safety-critical assets demand strong governance and vigilant monitoring for any changes. Threats Common pitfalls include overfitting in simulations, data leaks, missed constraint violations, fuzzy criticality, and apples-to-oranges metrics. To guard against these, use scenario stress tests, time-based backtesting, feasibility checks, clear tier definitions, and practical, comparable metrics.

IX. FUTURE RESEARCH AGENDA

Top priorities include building benchmarks that track changes over time, creating unified definitions for criticality, optimizing reliability and maintainability side by side, setting standard operational metrics, developing governance-first tools, and reporting surrogate errors transparently.

X. CONCLUSION

RBM optimization research is moving forward, but real-world readiness hinges on solid modeling of downtime and maintainability, realistic constraints, and transparent criticality tracking. The RAMS-C3 framework brings these elements together and provides a checklist to help teams compare and prepare for deployment.

11. This section dives deeper into the key themes that shape practical RBM optimization for RAMS, cost, and criticality. The aim is to turn the latest research (2020–2025) into actionable design choices, spotlighting where the evidence is solid and where assumptions are most common. While earlier sections covered the methods, here we explore how those methods are actually applied, which modeling decisions make the biggest difference, and which reporting details determine if others can use and replicate the work. we can use and repeat the work.

11.1 Theme A: "Risk" is not a single number—consequence modeling shapes decisions more than likelihood modeling.

Across the included literature, likelihood is often modeled with familiar reliability or degradation models: Weibull hazard, Markov transitions, competing risks, or state-dependent hazard driven by condition indicators. In many cases, different likelihood models produce similar relative rankings, especially when data are sparse, and uncertainty dominates. By contrast, consequence modeling frequently changes recommendations. If consequence is modeled purely as direct repair cost, preventive work tends to cluster around high-failure-frequency assets. If consequence includes service unavailability, preventive work shifts toward assets with long restoration times and high customer exposure. If safety is modeled as a constraint, the feasible plan region changes: certain deferrals become unacceptable regardless of cost efficiency.

The takeaway: RBM optimization should focus early modeling efforts on how consequences are structured. A simple hazard model paired with a detailed downtime and service disruption model can beat a fancy hazard model that overlooks downtime. This fits with RAMS thinking—availability depends on both how often things fail and how long they stay down. In practice, decision makers care more about the outages or harm they can avoid than about the technical hazard rate. That's why this review treats realistic consequence modeling as a top priority for readiness.

11.2 Theme B: Availability and maintenance. Most studies highlight gains in reliability, but few spell out improvements in availability—and when they do, the definitions jump around: instantaneous, operational, or service availability. Maintainability is even more inconsistently defined. Yet in the real world, maintenance plans are judged by how long outages last, how confident teams are in restoring service, and how predictable the workload can be and workload A RAMS model built for deployment should always split planned from unplanned downtime, since each is managed in its own way. Planned downtime can be scheduled with customer alerts and backup plans, while unplanned downtime is what really hurts a company's reputation. Maintainability should be broken down into actionable steps: detection, access, diagnosis, repair, testing, and getting back online. This breakdown lets teams see where investments will cut the worst-case downtime. Systematic reviews show that RAMS definitions and metrics are all over the map, making it hard to compare RBM optimizations. A simple fix—reporting both the average and the 90th percentile MTTR—can make results much clearer. interpretable.

11.3 Theme C: Criticality is a governance construct; tiered models are easier to adopt than weighted scores.

Many studies treat criticality as just a number that multiplies risk or downtime, which is easy for math but confusing in practice—a weight of 5 or 10 rarely means much to anyone. Different teams see criticality through their own lens: safety cares about hazards, operations about service, and finance about money. Industry 4.0 RBM says criticality should be updated as new evidence and conditions emerge, but if these changes aren't tracked, it can lead to confusion and shaky governance.

Tiered criticality offers a hands-on fix. Tiers set clear governance lines: Tier 1 assets need strict risk controls and top-level review, Tier 2 assets get service-level protection, and Tier 3 assets are open to opportunistic maintenance. Tiers can shift as redundancy or operations change, as long as every move is logged and reviewed. In optimization, tier-specific constraints produce results that stand up in audits: for example, 'Tier 1 safety risk stays below the threshold, Tier 2 downtime drops by X% within budget.' This approach matches ISO asset management's call for risk-based, objective-aligned planning.D: Constrained scheduling is where "optimization" meets reality. Even when studies develop elegant risk and cost models, execution often fails because resource constraints are not encoded. Constrained scheduling includes: outage windows, crew availability, skill constraints, precedence constraints, access permits, and spare parts lead times. In fleet contexts, hangar capacity and aircraft availability drive scheduling feasibility; in rail contexts, track possession windows and safety permits dominate. A dynamic scheduling framework for aircraft fleet CBM demonstrates how capacity constraints can force deferral and change the risk profile, even if the underlying condition prediction is accurate (Tseremoglou et al., 2024). A railway track geometry decision support model illustrates integration of cost, reliability, and availability into planning (Kasraei et al., 2022). Rolling stock policy development work further emphasizes that maintenance processes evolve and thus policies must accommodate change (Kumari et al., 2025).

When it comes to real-world adoption, realistic constraints matter more than which solver you pick. Mixed-integer programming is great for capturing every rule, but metaheuristics work just as well if they check feasibility and test sensitivity. For deployment, what counts most are three things: a schedule you can actually execute, a feasibility report covering crew hours and window compliance, and a clear picture of the residual risk after scheduling.

11.5 Theme E: Multi-objective optimization is most valuable when paired with a decision selection rule.

Many RBM problems are multi-objective: minimize risk, minimize cost, minimize downtime, and smooth workload. Multi-objective evolutionary algorithms are widely used because they produce Pareto sets. Yet, in practice, a Pareto set is not a decision. Governance requires a selection rule: a method for choosing one plan from the set based on thresholds, stakeholder preferences, or regulatory constraints. Surrogate-assisted multiobjective approaches enable efficient exploration when evaluations are expensive and are increasingly used in digital-twin-like scenarios practical way to choose among options is to use a tiered rule: set hard safety limits for Tier 1 assets, minimum availability for Tier 2, and then pick the plan with the lowest lifecycle cost that fits. Or, use a value-for-money rule—select the plan that delivers the most risk reduction per dollar, as long as the workload stays manageable. Without a clear selection rule, studies risk being seen as academic rather than truly useful for decision makers. her than operational decision tools.

11.6 Theme F: Dependency modeling shifts priorities and changes cost–availability tradeoffs. Assets rarely fail independently. Shared environments, common-cause failures, and functional dependencies cause correlated failures. Dependency also exists in maintenance execution: a single outage may enable multiple interventions, or a delayed spare may delay multiple tasks. A multi-objective availability and cost optimization study with subsystem failure dependencies demonstrates that dependencies change optimal solutions and the shape of the cost-availability frontier (Mellal & Zio, 2022).

For deployment, start simple with dependency modeling: group assets by redundancy or shared services. Even a rough map of dependencies can boost plan quality by avoiding underestimated consequences and spotting chances for joint work. To prevent misinterpretations during later reviews, authors should explicitly state in the abstract whether failures are assumed to be "independent" or "dependent." This small step can provide clarity and preemptively address potential misunderstandings.

11.7 Theme G: Evidence quality and inspection planning are part of the optimization problem. RBM is not only about maintenance actions; it is also about evidence acquisition. Inspection and monitoring decisions change uncertainty, and thus change optimal maintenance timing. Risk-based inspection planning provides a clear example: inspection timing can be optimized to reduce uncertainty and risk over the life of a deteriorating structure (Yang & Frangopol, 2021). The RBM literature can incorporate similar logic by treating inspection and monitoring as decisions with costs and expected value of information.

For deployment, organizations should see evidence programs—inspections, sensors, data pipelines—as investments to be prioritized with the same RAMS-C3 logic. If uncertainty is what drives overly cautious maintenance, then better evidence can cut both risk and cost.

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

11.8 Theme H: Learning-based methods require additional governance—drift, safety constraints, and explainability.

Deep reinforcement learning and similar methods promise adaptive policies for complex systems, but real-world acceptance hinges on monitoring for drift, ensuring safe exploration (which is often off-limits for safety-critical assets), and making results explainable. Even with strong simulation results, organizations will only deploy these methods if model updates are governed, validation is time-based, and residual risk is reported in familiar, practical terms.

Many studies leave out key details: vague cost assumptions, missing constraints, no time-based validation, and no sensitivity checks. Systematic review standards call for full transparency and completeness. RBM optimization research should follow suit by publishing exact objective functions (with units), clear constraint definitions, data split methods, parameter sources, and scenario setups. To encourage full transparency, a micro-checklist could be utilized as a prompt for authors: 'Have you shared the objective function units?', 'Have you clarified all constraints?', and 'Have you detailed your data split methods?' This concise micro-checklist aligns with PRISMA culture, nudging authors towards clearer reporting.

Many studies leave out key details: vague cost assumptions, missing constraints, no time-based validation, and no sensitivity checks. Systematic review standards call for full transparency and completeness. RBM optimization research should follow suit by publishing exact objective functions (with units), clear constraint definitions, data split methods, parameter sources, and scenario setups.

11.10 Theme J: A practical “minimum viable RBM optimization” blueprint.
Synthesizing the evidence, a minimum viable blueprint has the following components:

(1) Define criticality tiers and tolerability thresholds with governance ownership. Conduct this review quarterly to ensure alignment and accountability.

(2) Build downtime and MTTR distributions from historical work orders; separate planned/unplanned downtime. Update these distributions monthly for better accuracy and responsiveness.

(3) Create a simple risk model that maps failure likelihood to downtime and consequence by tier. Conduct an initial setup and review biannually, allowing adaptations as data improves.

(4) Implement constraint-aware scheduling with outage windows and crew capacity. Review and adjust this scheduling weekly to maintain effective execution with changing operational needs.

(5) Report residual risk by tier and produce an executable schedule. This should be part of a monthly operational review to keep task execution on track and transparent.

(6) Monitor execution outcomes and update downtime/cost parameters. A weekly cadence ensures you consistently capture deviations and quickly implement necessary corrections.

This blueprint can be This blueprint works even without advanced machine learning. It delivers quick wins by producing actionable work programs backed by clear risk logic. As data quality and governance improve, teams can add advanced methods like portfolio optimization, surrogate-assisted digital twins, and learning-based policies or measurements: what should be reported to compare studies.

To enable comparison across papers and domains, the review recommends reporting at least:
- Expected downtime hours avoided (availability metric).

- MTTR mean and tail (maintainability metric).

- Residual risk by criticality tier (governance metric).

- Lifecycle cost with clear decomposition (finance metric).

- Workload stability (execution metric).

These measures translaThese measures feed straight into operational planning and can be standardized to compare results across different domains.

The 2020–2025 literature shows meaningful progress in integrating optimization with risk modeling and in

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714089

expanding beyond single-objective reliability. The remaining adoption gap is primarily operational: downtime modeling, constraint realism, auditable criticality, and temporal validation. Addressing these gaps will enable RBM optimization methods to be evaluated and adopted based on implementability as well as mathematical novelty.

XII. EXTENDED METHODOLOGICAL GUIDANCE

12. Extended methodological guidance for review and implementation.

12.1 Review-method alignment with RBM topics.

Systematic reviews Engineering reviews often skip over key details like search methods, screening choices, and what data was pulled. PRISMA 2020 and PRISMA-S set the bar for transparency. For RBM optimization, this openness should cover not just literature steps but also modeling steps. Studies should spell out their objective functions (with units), constraint sets, decision timelines, and what information was available when decisions were made. These details show whether a planner could actually use the approach without hidden surprises. The fields that support synthesis.

This review suggests a minimum data extraction checklist for RBM optimization papers: domain and asset type, asset count, decision variables (like intervals or task start times), objective functions and weights, explicit RAMS outputs, how criticality is defined (static or dynamic), cost breakdowns (preventive, corrective, outage, inspection, spares), constraints (crew, outage, access, spares, precedence), how uncertainty is handled, evaluation design (synthetic, backtest, pilot), and reproducibility details (data access, pseudo-code, solver settings).

12.3 Example of a RAMS-C3 data extraction table.

A review paper can include a table that summarizes these fields for A moA review can include a summary table covering these fields for the most representative studies. Even if the full table is big, a compact sample boosts clarity. For instance, a table listing 10–15 key papers across different methods can show at a glance which ones modeled downtime, used tiered criticality, included outage windows, or applied temporal validation. This gives structured evidence without needing a full meta-analysis.

A practical deployment methodology works as an iterative spiral:

Phase 0—Scoping: define asset boundary, governance ownership, and what decisions the system will support (ranking, scheduling, or portfolio).

Phase 1—Baseline measurement: quantify current downtime, cost, and failure frequency by asset tier; establish data quality gaps.

Phase 2—Model development: build a simple risk model and criticality tiering; calibrate downtime distributions; validate on historical sequences.

Phase 3—Optimization and constraints: implement constraint-aware scheduling and portfolio rules; generate candidate plans; stress-test under disruption scenarios.
Phase 4—Integration: export work orders to CMMS/EAM; establish audit trails linking decisions to risk rationale; implement change control.

Phase 5—Continuous improvement: monitor outcomes, update parameters, and periodically review criticality tiers and tolerability thresholds (ISO, 2024a, 2024b).

12.5 Governance and assurance for safety-critical RBM.

For Tier 1 assets, governance often requires three additional practices. First, risk acceptance criteria must be explicit and approved. Second, the model must support traceability: which evidence led to which decision, and what residual risk remains. Third, change control must be formal: model updates must be versioned and reviewed, especially for data-driven elements. This is consistent with the general direction of asset management standards that emphasize risk-aware decision-making and continual improvement within an auditable system (ISO, 2024a, 2024b).

12.6 How to reduce “evaluation optimism” in RBM optimization studies.
Engineering optimization studies can inadvertently

overestimate improvements. Three practical steps reduce this:

(1) Use rolling-origin temporal evaluation when historical sequences exist.
(2) Introduce exogenous disruptions (crew absence, part delay, additional failures) as stress tests.
(3) Report feasibility rates and constraint violations for candidate schedules.
When learning-based methods are used, drift checks and decision-time information constraints should be documented (Zhang et al., 2020).

12.7 Summary.

This extended methods section builds on the thematic synthesis by laying out a structured approach for conducting and reporting RBM optimization reviews, and by turning the evidence into a step-by-step implementation method that fits governance needs and real-world operations.

REFERENCES

[1] El-Thalji, I., Al-Ghamdi, S. G., & others. (2025). Emerging Practices in Risk-Based Maintenance: Industry 4.0 Impact and the Next Generation of RBM. Applied Sciences, 15(3), 1159. https://doi.org/10.3390/app15031159

[2] Greiner, D., & Cacereño, A. (2024). Enhancing the maintenance strategy and cost in systems with surrogate-assisted multiobjective evolutionary algorithms. Developments in the Built Environment, 19, 100478. https://doi.org/10.1016/j.dibe.2024.100478

[3] ISO. (2024a). ISO 55000:2024 Asset management — Vocabulary, overview and principles. International Organization for Standardization.

[4] ISO. (2024b). ISO 55001:2024 Asset management — Asset management system — Requirements. International Organization for Standardization.

[5] Kasraei, A., & others. (2022). Maintenance Decision Support Model for Railway Track Geometry Maintenance Planning Using Cost, Reliability, and Availability Factors: A Case Study. Transportation Research Record, 2676(12), 1–14. https://doi.org/10.1177/03611981221077089

[6] Kumari, J., Karim, R., & Khanna, P. (2025). Dynamic maintenance policy development for railway rolling stock. Journal of Quality in Maintenance Engineering, 31(5), 23–42. https://doi.org/10.1108/JQME-01-2024-0005

[7] Mellal, M. A., & Zio, E. (2022). Multi-objective availability and cost optimization by PSO and COA for series-parallel systems with subsystem failure dependencies. Microprocessors and Microsystems, 89, 104422. https://doi.org/10.1016/j.micpro.2021.104422

[8] Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

[9] Pirbhulal, S., Gkioulos, V., & Katsikas, S. (2021). A Systematic Literature Review on RAMS analysis for critical infrastructures protection. International Journal of Critical Infrastructure Protection, 33, 100427. https://doi.org/10.1016/j.ijcip.2021.100427

[10] Rethlefsen, M. L., Kirtley, S., Waffenschmidt, S., et al. (2021). PRISMA-S: An extension to the PRISMA statement for reporting literature searches in systematic reviews. Systematic Reviews, 10, 39. https://doi.org/10.1186/s13643-020-01542-z

[11] Tseremoglou, E., Tsakalidis, A., & others. (2024). Condition-Based Maintenance scheduling of an aircraft fleet: A dynamic scheduling framework. Reliability Engineering & System Safety, 241, 109684. https://doi.org/10.1016/j.ress.2023.109684

[12] Yang, D. Y., & Frangopol, D. M. (2021). Risk-based inspection planning of deteriorating structures. Structure and Infrastructure Engineering, 18(1), 1–20. https://doi.org/10.1080/15732479.2021.1907600

[13] Zhang, N., Si, S., & Hu, C. (2020). Deep reinforcement learning for condition-based maintenance of multi-component systems with dependent competing risks. Reliability Engineering & System Safety, 205, 107084. https://doi.org/10.1016/j.ress.2020.107084